arXiv:1403.4054v1 [stat.CO] 17 Mar 2014

A general decision framework for structuring computation using Data Directional Scaling to process massive similarity matrices

Daniel John Lawson* and Niall Adams*†

March 18, 2014

Abstract

As datasets grow it becomes infeasible to process them completely with a desired model. For giant datasets, we frame the order in which computation is performed as a decision problem. The order is designed so that partial computations are of value and early stopping yields useful results. Our approach comprises two related tools: a decision framework to choose the order to perform computations, and an emulation framework to enable estimation of the unevaluated computations. The approach is applied to the problem of computing similarity matrices, for which the cost of computation grows quadratically with the number of objects. Reasoning about similarities before they are observed introduces difficulties as there is no natural space and hence comparisons are difficult. We solve this by introducing a computationally convenient form of multidimensional scaling we call 'data directional scaling'. High quality estimation is possible with massively reduced computation from the naive approach, and can be scaled to very large matrices. The approach is applied to the practical problem of assessing genetic similarity in population genetics. The use of statistical reasoning in decision making for large scale problems promises to be an important tool in applying statistical methodology to Big Data.

1 Introduction

If a dataset is too large to naively process with a desired model, we can either a) change the model, b) discard some data, or c) selectively order the calculation to yield useful partial results. This paper describes a framework for decision making to approximately evaluate complex calculations, with application to the specific

*Heilbronn Institute, School of Mathematics, University of Bristol
†Department of Mathematics, Imperial College London

problem of computing distance or more general similarity matrices. Statistical reasoning is used to navigate the computational challenges involved.

The framework has the character of a sequential decision problem. Specifically, at any sequential step a set of objects are selected for evaluation. The determination of which objects to evaluate is guided by a loss function. This loss function cannot be evaluated without the objects, and so we use an *emulator* to enable estimation of the loss. This emulator attempts to provide a computationally efficient prediction of unevaluated objects. To be of value, this computation must be less demanding than direct evaluation.

The general framework is applied to the problem of computing similarity matrices. There are numerous uses of similarity matrices in data analysis, including clustering, near neighbour search, and anomaly detection. In principle, evaluating the entire similarity matrix provides a solution for these uses. The current focus on "big data", fuelled by technology capable of automatically collecting and storing data at huge scale, provides new challenges for data analysis. Since the number of elements in a similarity matrix is proportional to the square of the number of objects, evaluating the entire matrix soon becomes computationally intractable. Additionally, we may not always require the complete evaluation of the similarity matrix. This paper develops a framework for a variety of data analysis activities based on similarity matrices that is intended to partially address the computational challenges of big data.

Data analytic uses of similarity matrices are used in numerous areas. Bioinformatics is one, with applications including sequence alignment Thompson *et al.* (1994), and population genetics Lawson and Falush (2012). A particularly topical driver is recommender systems (eg. Ricci *et al.* (2011)), which has motivated much work in the relatively new field of matrix completion (e.g. Candes and Plan (2010)). This field provides tools to impute the missing entries in a matrix. Whilst this is a very useful tool for computation, it is possible to make progress with less computationally intensive approaches.

The methodology we propose directly addresses scenarios involving data on N objects, where N is potentially very large. Objects need not be of the same dimension, but a similarity matrix S_{ij} can be constructed via a similarity function $S(i, j)$. A critical feature of these scenarios is that the similarities are available, but their number and computational cost make brute force evaluation intractable. This feature is an important contrast from the standard matrix completion scenario. Further, there is no feature space in which to place objects. Hence we use the evaluated objects to create a dynamic space via a process we call 'data directional scaling'.

Our framework uses statistical reasoning to carefully order a computation to obtain the maximum information benefit for the lowest possible computational cost. To do this, we draw on a number of statistical and machine-learning topics,

which are reviewed in Section 2. Our methodology is introduced in Section 3, including a simple (non-similarity based) example. We then define the similarity framework itself and introduce our computational framework, along with a simulation study to examine performance as we define the details. Section 4 examines how the approach scales with the size of the problem, with Section 5 providing a real-world application to a large genetics problem. We conclude in Section 6 with a discussion.

2 Literature review

This section reviews statistical, machine learning and computer science literature that is helpful or strongly related to our methodology. Section 2.1 discusses the use of emulators, Section 2.2 describes matrix completion, Section 2.3 describes efficient indexing strategies, Section 2.4 describes active learning, and Section 2.5 describes sequential decision making.

2.1 Emulation

Gaussian Process emulation Santner *et al.* (2003) is a major tool in the design and analysis of computer experiments, with applications ranging from climate models Rougier *et al.* (2009), and carbon budgets Kennedy *et al.* (2006) to demography Bijak *et al.* (2013) and testing cars in crashes Bayarri *et al.* (2009). In typical use, a limited number of computer experiments can be run and the goal is to choose the parameter values to run the expensive simulation model at, in order to fit a costly simulator to data. If the prediction space can be treated as continuous, then treating the observations as varying smoothly according to a multivariate Normal distribution proves surprisingly helpful in predicting good choices for where to evaluate the full simulator. There are naturally difficulties, one of which being that the cost of estimating a Gaussian Process grows rapidly with the size of the dataset. Decision making to exclude some data for future predictions Seo *et al.* (2000) can limit this when the data are added sequentially (e.g. Busby and Feraille (2008)). In this case diagnostics are of extra importance Bastos and O'Hagan (2009), to prevent the model from exploring the wrong area of the space.

Gramacy and Apley (2013) provide a computation-aware framework for Gaussian Processes. Whilst Gaussian Process emulators are by far the most common, other emulators are possible; for example, Šochman and Matas (2009) use an emulator over a binary space to speed computer vision. Bespoke emulators are also commonplace in the analysis of computer networks (e.g. Simmonds *et al.* (2000)).

Our work differs from the above in a) providing a general class of non-Gaussian Process emulator for similarity matrices; and b) extending the use of emulators

to the sort of tradeoff found in everyday computation in 'big data'. This is in contrast to the standard emulator framework in which the emulator has negligible computational cost.

2.2 Matrix Completion

The very active area of matrix completion is broadly concerned with imputing missing entries in a matrix from a relatively small proportion of observed entries. Candes and Plan (2010) provides a summary of recent work on matrix completion. The basic framework has two key features. First, the observed entries are assumed to be a random sample and are such that the observed matrix is full rank. Second, it is assumed that the full matrix has a low rank representation. Under these assumptions, algorithms implementing nuclear norm minimisation provide the computational means to attack the matrix completion problem. Notably proofs are available that a low rank solution can be recovered in few evaluations within a theoretical bound, in terms of the root-mean-square reconstruction error. A more recent review and a novel algorithm working under relaxed assumptions is provided by Hao *et al.* (2014). Indeed, many variations and refinements have been proposed, including Keshavan *et al.* (2010) which provides an algorithm that claims to open the door for matrix completion with big data. Scaling to big data, the target of our proposed framework, is a key challenge in matrix completion. To quote Candes and Plan (2010): "... but the computational challenges of solving problems with millions if not billions of unknowns obviously still require much research". The methodology and SOFT IMPUTE algorithm developed in Mazumder *et al.* (2010) provides capability to handle much bigger problems than many proposed matrix completion algorithms. Motivated by the SOFT IMPUTE algorithm and the fact the big data sometimes features sequential arrival, Bhojanapalli and Jain (2014) develop an online variant of matrix completion. They key aspect in this approach is a randomised technique for computing the singular value decomposition - a method that is at the core of all matrix completion algorithms.

In addition to the application of matrix completion to recommender systems discussed in the introduction, other interesting applications include genotype imputation Chi *et al.* (2013), functional genomics Gerlee *et al.* (2013), computer network performance prediction Liao (2013) and global positioning in sensor networks (eg. Biswas *et al.* (2006)). This is an interesting example because there are often power constraints on sensors that means some distances are not observed, leading to a bias in the sampling mechanism. This is explored in Taghizadeh (2014).

Matrix completion would be a valid approach to some of the problems we face in this work. However, it is not the approach we take for several reasons. First, we want to make informed decisions with little information, when a matrix

completion would be of insufficient rank. Second, the requirement to obtain a full rank solution also provides a restriction on which elements can be evaluated, as all rows and all columns must be visited. Third, little attention has been paid to how to update a matrix completion as data is observed sequentially. Fourth, we can make use of limited information without performing a full matrix completion. Finally, the literature does not provide an 'off the shelf' solution tailored to the particular type of similarity matrices we face; as we shall see below, the current methods perform badly because they make the wrong assumptions about the form of the matrix.

2.3 Indexing methods

There are numerous data structures useful for indexing multivariate spaces, including KD-trees Bentley (1975), Quadtrees Finkel and Bentley (1974) R trees Guttman (1984), and X-trees Berchtold *et al.* (2001). Our problem does not observe a multivariate space, but instead we view similarities as lying on some implicit manifold in an *unknown* space. These approaches therefore cannot be directly applied. However, if we use some other algorithm to recover the space then indexing can be of value to allow rapid lookup of neighbours. To do this the tree algorithm would need to handle a growing space as the algorithm progresses, for which we are not aware of any easily applicable framework.

2.4 Active learning

Some supervised classification problems have the following characteristics. First, there is a small amount of labelled data and a large amount of unlabelled data. Second, the process of labelling is costly. In these scenarios, active learning Cohn *et al.* (1996) is used to select unlabelled data for presentation to an oracle for labelling. The objective is select those unlabelled data that will yield the greatest improvement to the classification model, thereby minimising the cost and maximising the utility of labelling. There are numerous variations on this basic theme. Settles (2010) provides an excellent review, and details the many heuristics that have been explored. Interestingly, it seems only recently that performance comparison against a random selection benchmark is important.

Our framework shares many characteristics with Active Learning: we seek to score unevaluated similarities S_{ij} with the intention that the scores are indicative of the value of the similarity to the task. Candes and Plan (2010) discusses the role of an oracle in matrix completion problems. As with active learning, random selection of similarities is an important benchmark in our framework. Unlike most work in Active Learning, we obtain a continuous outcome, we do not have an explicit feature space in which to work, and we are interested in different loss functions.

2.5 Sequential decision making

Sequential decision often involves selecting actions to maximise reward in noisy or uncertain environments. A good overview is given in Cesa-Bianchi and Lugosi (2006). The "exploitation-exploration dilemma" is always present in such problems. This dilemma concerns the competing objectives of selecting actions for maximal gain (exploitation) and for reducing uncertainty (exploration). The core of our framework is sequential selection of objects S_i to evaluate, and the exploitation-exploration dilemma naturally arises.

One simple approach to address the dilemma is to incorporate an element of randomised decision making. Such methods, called ε-greedy methods, select the greedy action (the action with highest predicted reward) with probability $1 - \varepsilon$, and a random non-greedy action with probability ε. Such algorithms have solid theoretical justification for infinite horizon problems and have been demonstrated to behave well empirically.

3 Methodology

We start in Section 3.1 with the general decision framework for ordering computation. Section 3.1.3 illustrates the advantage of reasoning about computation with the well known example of an autoregressive model. Section 3.2 explains the Data Directional Scaling method for constructing a similarity space. Section 3.3 defines the similarity problem precisely, whilst 3.4 fully defines the emulation framework. Section 3.5 defines the choice framework, whilst the simpler issues of Assessment (Section 3.6) and Termination (Section 3.7). Finally, this is wrapped up with a simulation study in Section 3.8.

3.1 General framework

Consider a discrete set of objects $i = 1 \cdots M$ about which we can choose to take measurements (hereafter called computations) $S_{i=1 \cdots M} = \{S\}$ at cost L_i, with mean cost $\mathbb{E}(L_i) = L$. If ML is large relative to our computational budget, we cannot afford to compute them all. Despite this, we wish to compute some quantity, say $\phi(\{S\})$, of the whole dataset. How well we estimate ϕ is defined by a loss $\mathcal{L}(\{S\})$. For example, if $\phi = \bar{S}$ is the sample mean then we could define $\mathcal{L}(\{S\}) = \text{var}(\phi)$ to specify a minimum variance solution.

We will iteratively make choices about which object $i(t)$ to evaluate at iteration $t = 1 \ldots T_{max}$. The number of iterations T_{max} may not be fixed. Let $\Omega(t) = \{i(1), \ldots, i(t)\}$ be the set of objects selected up to now, and $S^*(t) = \{S_{\Omega(t)}\}$ be the associated observed quantities. Our framework uses two key concepts:

- **A Choice function**: $\mathcal{C}(S^*(t))$ takes the previously observed information $S^*(t)$ and chooses an object i from the unobserved objects $\Omega^{\perp}(t)$.

- **An Emulator**: $\mathcal{E}(i; S^*(t))$ The emulator takes the previously observed information $S^*(t)$ and returns a predictive distribution $\hat{S}_i(t)$ on any observed object i.

The choice function $\mathcal{C}(S^*(t))$ can use any of the available information. Specifically, the choice function can use the emulator $\mathcal{E}(i; S^*(t))$ to make intelligent decisions via the following tools:

- **A Loss estimator**: $\hat{\mathcal{L}}(S^*(t)) = \mathbb{E}(L(\hat{S}))$ over possible values of $\hat{S}_i(t)$. Note that $\hat{S}_i \equiv S_i$ for $i \in \Omega$.
- **A Heuristic**: $\mathcal{H}(S^*(t))$ is a decision rule avoiding emulation. In some special cases, it can be shown to minimise a loss $\mathcal{H}(S^*(t)) = \arg\min \hat{\mathcal{L}}(S^*(t))$. If it does, it is *exact* for that loss.

Although emulators that make a point prediction could be exploited, most interesting loss functions require as assessment of uncertainty. We will explore the loss function more fully in Section 3.5 after introducing the similarity problem.

3.1.1 Computational constraints limit choice

The purpose of this framework is to reduce computation. As this restricts the number of quantities that the Choice function can evaluate, it is helpful to describe it in more detail. $\mathcal{C}(S^*(t))$ proposes a number $m(S^*(t), t)$ of objects $\Omega_m^{(t+1)}$ for which the minimum predicted loss $\arg\min_m \hat{\mathcal{L}}(S^*(t)|\Omega_m^{(t+1)})$ is estimated. A special case is where there is only one proposal, in which case no loss calculation is required. This can happen if an exact heuristic $\mathcal{H}(S^*(t))$ is available for the loss $\mathcal{L}$. Even if there is no known exact heuristic, proposals from other heuristics are often helpful.

Finally, note that the Choice function has available the current iteration t as well as being able to assess the success of previous choices. This can be used choose to enable or disable proposals. For example, it is often useful to use more complex proposals initially, and then switch to simpler proposals as the marginal return of careful choice decreases.

3.1.2 Full choice algorithm

The algorithm proceeds as follows. For each iteration t:

1. *Choice*:
   (a) *Proposal*: $\mathcal{C}(S^*(t))$ proposes objects $\Omega_m^{(t+1)}$ to evaluate.
   (b) *Emulation*: If $M > 1$, estimate the loss $\hat{\mathcal{L}}_m(\hat{S}(t+1)|\Omega_m^{(t+1)})$.
   (c) *Decision*: Choose the action m' that minimises the loss.

2. *Evaluation and Assessment*: Compute $S^*(t+1)$ of the chosen objects $\Omega_{m'}^{(t+1)}$. Compare the predictive distribution $\hat{S}_{\Omega_{m'}^{(t+1)}}(t)$ to the observations $S^*_{\Omega_{m'}^{(t+1)}}$.

3. *Termination*: Stop if a stopping criterion is reached.

These processes are precisely defined below. Specifically, Section 3.4 deals with emulation, which also uses the results of assessment. Section 3.5 deals with the remaining aspects of the Choice function. Section 3.7 addresses stopping the algorithm.

3.1.3 Simple example

Before we narrow our attention to similarity matrices, it is helpful to demonstrate the value of reasoning about computation more generally. As an example, let S_i (with $i = 1 \cdots M$) describe an autoregressive (AR1) model:

$$S_i = \psi + \phi S_{i-1} + \epsilon_i$$

with $\phi < 1$ and $\epsilon_i \sim N(0, \sigma_\epsilon^2)$. Let us assume that the process has reached stationarity. We wish to estimate $\mathbb{E}(S) = \mu = \psi/(1-\phi)$ via an estimator $\hat{\mu}$. If either M is large, or the cost per computation L limits the number we can make, then we cannot compute $\mathbb{E}(S)$ using $\bar{S} = \sum_{i=1}^{M} S_i$. However, we can obtain a Monte-Carlo estimate by evaluating a subset $T_{max} < M$. Since they are correlated, which should we choose?

Let $\hat{\mu} = \sum_{i \in \Omega(T_{max})} S_i$, and the loss be $\mathcal{L}(S) = \mathrm{var}(\hat{\mu})$. From the standard Monte-Carlo methodology (e.g. Gamerman (1997)) it is known that minimising this variance is the same as minimising the covariance between the samples. This can be performed explicitly to result in an optimal heuristic.

If T_{max} is known in advance then $\mathcal{H}(S^*(0))$ is known, i.e. we can choose $\Omega(T_{max})$ before we compute any S. This uses evenly spaced samples with separation N/T_{max} (called thinning in the Markov-Chain Monte Carlo literature). If however T_{max} is not known in advance, $\mathcal{H}(S^*(t))$ can use an iterative greedy approach by picking the object $i(t+1)$ furthest from all previous objects $\Omega(t)$, i.e. select objects $1, N, N/2, N/4, 3N/4$ etc. If the autocorrelation length of S is greater than T_{max} then this approach massively reduces the variance over the naive approach of using the first T_{max} objects. It also clearly improves on the selection of random objects, which is a heuristic that is frequently helpful when nothing else can be done.

If our desired loss function is to instead estimate the parameters ψ and ϕ as accurately as possible, then we must decide on an estimation framework (least squares, Yule-Walker equations, maximum likelihood via numerical optimization, etc. Hamilton (1994)). Then we face a tradeoff between obtaining high lags to estimate μ and low lags to estimate ϕ. In this case, the correct model can be

used via a Kalman filter KALMAN (1960). This can be seen as Emulation, and is one way to handle missing data HOWELL (2007); JONES (1980). Developing the appropriate loss and choice functions is possible but outside the scope of this paper.

3.2 Data Directional Scaling

Before we go into the specifics of decision making for similarity matrices, we introduce 'Data Directional Scaling' (DDS) which is a type of Multidimensional Scaling (MDS YOUNG and HAMER (1987)) that has proven very helpful for our emulator. In MDS, the data is seen as having a position on a manifold or metric space described by a basis. There are a wide range of MDS methods, but the most commonly used is principal component analysis (PCA or equivalently, singular value decomposition, SVD). In PCA, basis vectors are chosen to be orthogonal and decreasing in their contribution to the variance observed in the dataset.

For the similarity problem, consider a set of objects $i = 1 \dots N$ describing data D_{il} with $l = 1 \cdots L_i$ indexing information about those objects. There is not necessarily a consistent dimension to the data, but a non-negative similarity function $S_{ij} = S(\mathbf{D}_i, \mathbf{D}_j)$ can always be evaluated. If additionally S forms a metric space then $S_{ii} = S_0$ for all i and we can write distances $X_{ij} = S_0 - S_{ij}$. X_{ij} are also non-negative with $X_{ij} = 0$ if and only if objects i and j are identical under S. Conversely, if S is non-metric, we can still convert to a distance-like measure via $X_{ij} = \max_k S_{kk} - S_{ij}$ but some self-distances are non-zero. We can interchangeably work with S or X as we do not exploit any properties of a metric.

Without access to any space in which the objects can be embedded, most traditional inference frameworks are difficult to apply. The matrix completion approach seeks to find the underlying space by inferring a low-rank latent space (typically an SVD). However, this fails when insufficient matrix elements have been evaluated.

Our approach evaluates matrix elements in rows $\mathbf{S}_i$ by calculating all similarities with object i. This is called 'evaluating object i'. It leads directly to a vector space $\mathcal{S}(t)$ using only the information from observing objects $\Omega(t)$. The basis of $\mathcal{S}(t)$ is the *columns* $\mathbf{S}^*_b$ of the similarity of each evaluated object $k \in \Omega(t)$ with each other object $b \in \Omega(t)$. Therefore an additional observation at iteration $t+1$ simply creates a new direction in the space. Figure 1a illustrates this procedure.

This is useful because columns $\mathbf{S}^*_i$ for $i \in \Omega^\perp(t)$ are observed. The emulator $\mathcal{E}(i, S^*(t)$ provides a position $\alpha_i(t)$ for object i in the space $\mathcal{S}(t)$. For practical reasons, we define $\alpha_i(t)$ to live in the $t-1$ dimensional simplex; i.e. $\sum_{\alpha_i} = 1$ and $\alpha_i > 0$ for all i.

Now our objects i are standard observations in a vector space, and we can use standard tools to perform inference. Specifically, we treat columns $S^*_i(t)$ as observations, and rows $S_i(t)$ as responses. We can therefore use standard linear

methods to predict unobserved rows $S_i(t)$. This procedure is discussed in Section 3.4.

Figure 1b interprets this procedure when objects are truly embedded in $\mathbb{R}^2$, which is a metric space and therefore the similarities are symmetric. If we further assume that similarities are observed with noise, then triangulation is non-trivial. By observing objects at the 'corners' of the space, the unobserved objects fall inside or near the convex hull of the observed objects. In this case, we will be able to accurately represent the whole space without extrapolation.

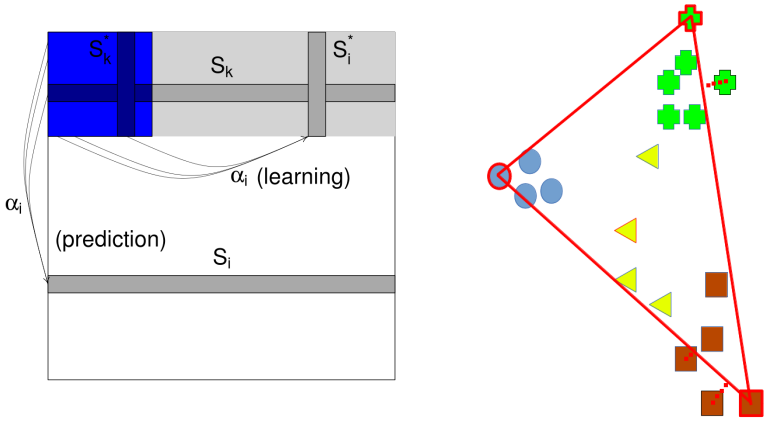

Figure 1: Data Directional Scaling illustration. Left: Quantities used in DDS. Observed rows of S are shown at the top, with the vector space $\mathcal{S}(t)$ being defined by the set of columns of observed similarities S^*_k for observed objects $k \in \Omega t$. Object i is associated with observations S^*_i, which is mapped to a position α_i in the vector space using the learning model for the Emulator $\mathcal{E}(i; S^*(t))$. The prediction space is the rows S_i of the matrix, for which we observe S_k with $k \in \Omega(t)$. Right: Example with an underlying Euclidean space $\mathbb{R}^2$. If we have observed the three objects in red, then we can reconstruct any point in the $2D$ simplex (i.e. triangle) shown in red. Positions outside the simplex are mapped to the nearest point inside the simplex.

3.3 Similarity model

We will use DDS and the choice framework to optimally select elements from the similarity matrix to evaluate. The expected cost of a similarity computation is defined to be L, with $L \propto \mathbb{E}(L_i)$ if the computation is linear. The total cost of computing this matrix is therefore $O(N^2 L)$, which we assume is too large. Our methodology will reduce this computation to $O(N^2 + NL)$ for quite general problems. This can be reduced to $O(NL)$ if a) the full matrix is not required

and b) either a less exact solution is satisfactory, or explicit solutions to certain sub-problems (i.e. exact heuristics) can be found.

As above, we use a choice function and an emulator. The Choice function $\mathcal{C}(S^*(t))$ now can return either a single similarity element (i.e. an i, j pair), or a set of elements. To use DDS, it is necessary to evaluate S_{ij} in entire rows. The emulator still returns the value of single elements, and takes the form

$$\mathcal{E}(i, j; S(t)) = \hat{S}_{ij}(t) \sim N(\bar{S}_{ij}(t), \sigma^2_{ij}(t)) \tag{1}$$

which provides a distribution for any element i, j. $\bar{S}_{ij}(t)$ and $\sigma^2_{ij}(t)$ both require further specification. Whilst it is arguable whether the Normal model is appropriate, we note that it is justifiable on computational grounds. Further, if the similarities are sums over features then (as in Gaussian Process emulation) the Central Limit Theorem justifies this choice.

3.4 Emulation

Recall that in our framework, the utility of an emulator is to be able to predict how calculating a particular set of elements will improve our ability to optimise a loss. We therefore will need to be able to compute many emulated values, and recompute our emulator as new information arises. A slightly separate use is to 'complete' the matrix at the end of the algorithm to provide a 'best estimate'. Compare this use with matrix completion, for which the computational budget is allowed to be larger.

For defining the emulator, we abuse our notation and set $\Omega(t)$ to be the index of the objects we have evaluated, i.e. *full rows* of the matrix S evaluated up to time t. The emulator takes the form:

$$\hat{S}_{ij} \sim N\left(\bar{S}_{ij}(\Omega(t)), \delta^2_i(\Omega(t))\right), \tag{2}$$

where we have dropped the obvious dependence on S^*. The quantities δ_i and $\bar{S}_{ij}$ are both inferred using linear prediction frameworks, due to computational restrictions.

3.4.1 Prediction of the mean

Using the Data Directional Scaling from Section 3.2 and Figure 1a, $\bar{S}_{ij}(\Omega(t))$ is estimated as follows. For $k \in \Omega(t)$, a linear regression is used to *learn* the observations in column i using the observed columns k:

$$S^*_{ji, j\in\Omega(t)} = \sum_{k\in\Omega(t)} \alpha_{ik} S^*_{jk} + e_{ji}(t).$$

Row i is *predicted* using the observed rows:

$$\bar{S}_{ij,j\in 1\cdots N} = \sum_{k\in\Omega(t)} \alpha_{ik} S_{ij} + d_{ij}(t).$$

Here, $e_{ji}(t) \sim N(0, \epsilon_i^2(t))$ is the standard residual error on observed elements $S^*(t)$. $d_{ij}(t) \sim N(0, \delta_i^2(t))$ is the out-of-sample prediction error for an unobserved element S_{ij}. The residual $\epsilon_i^2(t)$ is useful, but we will are particularly interested in the understanding the 'emulator variance' $\delta_i^2(t)$, the same quantity from Equation 2.

As discussed in the DDS section, α_i is learnt under the constraint that $\sum_k \alpha_{ik} = 1$ and all $\alpha_{ik} \geq 0$ using one of the following estimation procedures:

NN: Nearest Neighbour. Choose $\alpha_i = \arg\min_\alpha(\| \sum_{k\in\Omega(t)} \alpha_{ik} S_{jk} - S_{ji}\|_\infty)$. By weighting large similarities 'infinitely' strongly this norm leads to selecting $\alpha_{ik} = 1$ for $k = \arg\min_j(S_{ji})$, i.e. the nearest neighbour in the column S_{ji} with $j \in \Omega(t)$, and zero otherwise.

MM: Mixture Model. Choose $\alpha_i = \arg\min_\alpha(\| \sum_{k\in\Omega(t)} \alpha_{ik} S_{jk} - S_{ji}\|\|_2)$, i.e. fit the mixture to minimise the square error. This is inferred using quadratic programming Turlach and Weingessel (2013).

Figure 1b captures the mixture model. We also considered an unconstrained linear regression, but this is unstable because the inference is overspecified when $t > R$; in practice two objects within one cluster are used with a very large positive and negative weight.

Because we are dealing with similarities, it is necessary to handle 'self' specially. We swap S_{kk} and S_{ki} so that the estimator 'sees' the similarity of k with i instead of the that of k to k (which is only zero for metric spaces).

3.4.2 Prediction of the Emulator Variance

Estimating δ_i must be rapid, and capture the distribution of the sampled objects and structure of the manifold on which they lie. In principle it is a function of the full set of similarities with the evaluated objects, and is not 'observable' in the way $\bar{S}_{ij}$ was, because the uncertainty of an observed element is zero. We calculate the uncertainty of the prediction via a linear combination of most conceivable summary statistics:

$$\delta_i^{\text{full}}(\Omega(t)) = \gamma_0 + \gamma_t t + \gamma_{t^{-1}} t^{-1} + \gamma_\epsilon \epsilon_i(t) + \sum_{n\in\mathcal{N}} \gamma_n \bar{R}_n(i; \Omega(t)) + \text{residuals}. \quad (3)$$

This includes coefficients for an intercept γ_0, the number of previously evaluated objects t via (t and t^{-1}), and the regression residuals $\epsilon_i(t)$. We also account for

the full similarities distribution of object i with evaluated objects $k \in \Omega(t)$ using 'similarity distance' in the L_n norm (using $n \in \mathcal{N} = \{1, 2, \infty\}$):

$$\bar{R}_n(i; \Omega(t)) = \left(\sum_{i \in \Omega(t)} \|S_{ji}\|^p \right)^{1/n}.$$

Fitting the emulator variance $\delta_i(t)$: The set of objects $\Omega(t)$ chosen for evaluation may be very far from random and S_{jk} (for $j, k \in \Omega(t)$) are typically biased to small values (i.e. the distances between them are larger than average, as in Figure 1b). Extrapolation using linear regression can therefore be misleading and we enforce a sensible prediction by using non-negative least squares regression (NNLS,Lawson and Hanson (1974)) instead. This automatically sets some coefficients to zero, although explicit penalisation can also be used.

Empirically, our NNLS metho most commonly chooses the regression:

$$\delta_i(\Omega(t)) = \gamma_\epsilon \epsilon_i(t) + \gamma_\infty \bar{R}_\infty(i; \Omega(t)) + \text{residuals}. \tag{4}$$

This is computationally efficient to work with as only nearest neighbours have a changed γ_∞ coefficient and $\epsilon_i(\Omega(t))$ changes only by a single coordinate, leading to an efficient calculation for $\delta_i(t+1)|\delta_i(t)$. Therefore the results we report use Equation 4, although we have implemented Equation 3 and checked that its performance is not significantly different.

Cross validation estimate of $\delta_i(t)$: We have a choice for how we fit $\delta_i(t)$ to achieve the best predictive power for the lowest computational effort. We have available to us, for free, the observed values of $\delta_j(\Omega(t'))$ for the previously chosen objects $j \in \Omega(t)$ for previous times $t' \leq t$ with a different set of observed objects $\Omega(t')$. This acts as a 'poor mans cross validation'. We can either:

1. Use only the observed history;
2. Use only cross-validated estimates from the current observations t, by treating each object $i \in \Omega(t)$ as if it were the last object to be evaluated;
3. Combine the approaches by retaining any cross-validation performed at previous iterations and combining it appropriately with cross-validation at iteration t.

Using only the observed history will cause problems when there are few evaluated objects. Additionally we might expect 'non-stationarity' in that early objects are measured in a different space and hence follow a rather different distribution to later objects. Conversely, using only cross-validation at the current iteration is computationally costly, does not account for any time-dependent learning, and exposes only a limited set of distances. Combining the approaches would seem to

be appropriate but requires choosing a way to 'age-off' less accurate information from early in the process. Additionally, we might wish to choose how much cross-validation to perform; intuitively, less is needed later in the process.

We therefore define:

- **A cross-validation operator** $\mathcal{V}(t)$ which decides which historical objects to retain.

By default we use $\mathcal{V}_{obs}(t)$ which uses all the history up to time t (excluding the first 5).

We have an additional choice about how to handle the additional information about S_{ik} when $k \in \Omega$ if $\mathcal{S}$ is known to be symmetric. First, we can set $\hat{S}_{ik} = S_{ki}$ and use the above framework. Second, we can use its inferred value to calibrate $\delta_i(t)$, estimating $\delta_i(t) = \sum_{k \in \Omega(t)} (\hat{S}_{ik} - S_{ki})/t$ instead of the regression described above. We have not investigated these options in detail, instead treating all similarities as non-symmetric.

3.5 Choice: Proposals and Decisions

We can exploit the emulator as defined to make intelligent decisions about which objects and/or matrix elements to evaluate using the choice function $\mathcal{C}$. The loss depends on the emulated matrix $\hat{S}(t)$ via the evaluated similarity elements $S^*(t)$. It also depends on the computation spent. Let L_{ij} be the computational cost for computing element S_{ij}. Then for simplicity we consider losses of the form

$$\mathcal{L}(S^*) = \mathcal{L}_0(S^*(t)) + \mathcal{L}_1 \left(\sum_{i,j \in \Omega(t)} L_{ij} \right),$$

where $\mathcal{L}_0$ is the 'model fit' loss, and $\mathcal{L}_1$ is the 'computational cost' loss. This simplification means that we can minimise each loss separately and easily combine them. If we further assume that $L_{ij} = L$ for all pairs, then the computational cost of any proposal containing the same number of elements is the same. For notational simplicity we therefore omit the subscript $\mathcal{L}_0$ below.

Loss for model fit: A fairly general loss is a (weighted) mean prediction error of the form:

$$\mathcal{L} = \mathbb{E} \left(\sum_{ij} w_{ij}(\hat{S}_{ij}) \| \hat{S}_{ij} - S_{ij} \|_n \right), \tag{5}$$

where the weights $w_{ij}(\hat{S}_{ij})$ allow us to specify which elements are most important for a particular problem, and can allow for covariance. We consider the L_n norm of the difference between the true and observed values of the matrix.

We do not have access to the true value of S_{ij} at any iteration, but the emulator provides a calibrated way to estimate the expectation for decisions. Therefore,

although we do not have access to the absolute value of the loss, we can estimate the difference in loss between decisions with high accuracy, and hence make good decisions.

It is helpful to distinguish between two classes of evaluation that we use:

- **Global search:** To identify bulk cluster structure, propose to evaluate an entire row of S, setting $\Omega_m^{(t+1)} = i$. Evaluate all similarities S_{ij} of an object i with all other objects j.
- **Local search:** To identify local neighbourhood structure, propose to evaluate a single element i, j with value S_{ij}. This is given the notation $\Omega_m^t = i, j$, and if chosen, $\Omega(t)$ now includes this element. i.e. t is *not* incremented.

3.5.1 Global Search

If we were to evaluate all similarities with object i, then $\Omega_m^{(t+1)} = i$. This has two effects on the loss. First, it evaluates $\hat{S}_{ij}$ to be S_{ij}, and second, it can be used to improve predictions of other objects via the emulator. Evaluating the improvement of the loss in principle requires integration over all assignments of $\hat{S}_{ij}$.

However, we can obtain a good estimate of $\hat{\mathcal{L}}_m(\hat{S}(t+1)|\Omega_m^{(t+1)})$ for significantly less compute by 'plugging-in' the expected value $\bar{S}_{ij}$. We therefore assume that a given row will be the mean predicted value, and predict the change in loss for the rest of the matrix. Therefore we can write the loss and therefore the optimum choice $\mathcal{C}(t+1)$ as

$$\hat{\mathcal{L}}_m(\hat{S}(t)) = \sum_{ij} w_{ij}(\bar{S}_{ij}) \|\delta_i(t)\|_n \tag{6}$$

$$\mathcal{C}(t+1) = \arg\min_m \left(\hat{\mathcal{L}}(\hat{S}(t)) - \hat{\mathcal{L}}_m(\hat{S}(t+1|\Omega_m^{(t+1)}) \right) \tag{7}$$

where $\delta_i(t)$ for observed rows is defined to be zero. Notice that this loss depends directly on the emulator variance only. It depends on the emulator predictions via their interaction with the variance (and potentially via the weights). The expectation in Equation 7 is easily evaluated by estimating $\delta_j(t+1|\Omega_m^{(t+1)})$ via the same regression from Equation 4, with the set of objects $\Omega_m^{(t+1)}$.

We define one loss and two heuristics via this framework.

- *RMSE loss*: The l_2 norm is a natural measure and can be used to directly minimise the root-mean-square error (RMSE) in the prediction. This is a matrix completion approach. This loss can be computed by brute force in $O(N^2)$ by evaluating the effect of each observation i on the corresponding variance δ_j^2.

- *Furthest distance heuristic*: If we use uniform weights and the l_∞ norm for the loss, then the largest value of δ_i dominates. The loss is minimized by evaluating the object that is furthest away from all evaluated objects. This heuristic is inspired by the nature of the space from Section 3.2.
- *Random selection heuristic*: Whilst not strictly available under the loss framework, if all weights are zero the loss is uniform and we can select at random. This is an important benchmark.

In practice, evaluating the full emulated RMSE loss function is too costly. It is useful instead to use a Monte-Carlo estimate of it. We estimate the loss using a constant number $p = \min(80, N - t)$ of random unevaluated objects (with the addition of one chosen by the furthest distance heuristic). We then estimate the loss using those samples only, reducing the computation from $O(N^2)$ to $O(p^2)$.

3.5.2 Local Search

Local search permits evaluation of elements without calculating an entire row. However, our emulation framework cannot exploit this information without additional complication. We therefore allow local search to only influence the loss at specific values of i, j and omit any learning that this point provides about the rest of the matrix. For some loss functions, however, it is extremely important to calculate certain matrix elements at the omission of others and therefore these choices might be applicable. We can choose local proposals using the same loss function as in the global case, and indeed decide between choices on the basis of this loss.

In this case we need to consider $\mathcal{L}_1$ in the loss. In practice we have used a linear function $\mathcal{L}_1(x) = cx$. In this case the cost of a local proposal is cL and the cost of a global proposal is cLN. The factor c is chosen to convert computational time units to 'loss per iteration' units.

3.6 Evaluation and Assessment

Evaluation and assessment are straightforward. The matrix elements are evaluated, and compared to the predictive distribution. The predicted mean of the emulator $\bar{S}_{ij}(t)$ is compared to the observed values S_{ij} using its RMSE:

$$\text{RMSE}(\bar{S}_{ij}(t)) = \sqrt{\frac{1}{N}\sum_{j=1}^{N}\left(\bar{S}_{ij}(t) - S_{ij}\right)^2}.$$

The predicted emulator variance $\delta_i^2(t)$ of the emulator is also compared to the empirical residual variance $\delta_{obs,i}^2 = (1/N)\sum_{j=1}^{N}\left(\bar{S}_{ij}(t) - S_{ij}\right)^2$ using its error:

$$\text{Error}(\delta_i(t)) = (\delta_i(t) - \delta_{obs,i}(t))^2.$$

This is much noisier since we only obtain one error per object observed.

In both cases, it is of value to determine convergence which could be used in the termination criteria. There are standard test statistics that could be used. We have not tried to terminate the algorithm on this basis because the loss functions we have used do not converge, but instead slow down. Other loss functions may converge, however.

3.7 Termination

There are three fundamental ways of running the algorithm, which can be defined using an appropriate computational loss term $\mathcal{L}_1$. Termination can be seen as an option that minimises the loss $\mathcal{L}$, when all possible actions result in a decrease in loss.

1. Obtain the minimal loss for a fixed computation. We have seen in Section 3.1.3 that it is helpful to think about the total number of evaluations that will be available. However, most calculations are not this structured and we can run our algorithm until a computational budget is reached. The corresponding computational loss $\mathcal{L}_1$ takes an arbitrary function up to the budget B when it is set to ∞.

2. Obtain the cheapest estimate with a fixed precision. If a desired precision of the loss is known *apriori*, then we can use standard results from sequential stopping rule theory Lindley (1961); Berger (1985) under the assumption that the loss has normally distributed error. Again, $\mathcal{L}_1$ takes a step to ∞ but now when the precision is reached.

3. Use the full loss $\mathcal{L}_1$. We would then be able to explore the cost/benefit ratio of obtaining further similarities by making further computations.

We have here only worked with criterion 1, i.e. we compare the performance for a given amount of compute. Further, in the simulated example, the cost of each computation L is arbitrary and so we do not compute the real computation performed but estimate it to a factor of L.

3.8 Simulation Example

We illustrate our methodology with application to a simulated dataset. $L = 1000$ Gaussian features were simulated by assuming a tree correlation structure for $N = 500$ samples in $R = 10$ evenly sized clusters. This type of dataset is representative of many real world applications, including genetics. It is also challenging for standard matrix completion approaches because the eigenvalues correspond to the clusters Lawson and Falush (2012), which might not be informative about all

samples. This assumption is required for matrix completion to have performance guarantees Candes and Plan (2010). Figure 2 illustrates one such matrix.

To assess the robustness and performance characteristics of the algorithm, we vary the difficulty of the problem. For this we vary the correlation within and between clusters, and allow outliers of varying sizes. Table 1 describes the parameters used for the datasets we generated, which lie on a continuum from 'clustered' in which inference is simple, to 'corrupted' which contains many outliers and has weaker correlation structure.

Parameter	Description	Start 'clustered'	End 'corrupted'
c_0	Within cluster correlation	0.75	0.3
c_1	Close cluster correlation	0.5	0.25
c_2	Distant cluster correlation	0.25	0.15
a_h	Inverse Outlier α parameter	0	0.2
b_h	Inverse Outlier β parameter	0	5

Table 1: Simulation dataset parameters for the correlation between samples. There are 10 clusters, with close clusters being (1-2,3-4,5-6,7-8,9-10), and distant clusters are (1-4,5-8). Outliers are generated by mixing rows with weight $1 - w_i$ with a 'self' direction (1 on the diagonal, 0 off it). The weights $w_i = \text{beta}(\alpha = 1/a_h, \beta = 1/b_h)$ are biased to 1 and are exactly 1 when $a_h = b_h = 0$. 11 datasets are generated, evenly space from the 'clustered' to 'corrupted' parameters (referred to as a difficulty scale).

The 'clustered' dataset contains distinct clusters, whilst in the 'corrupted' dataset, clusters overlap and most samples are subject to significant deviations from their cluster. The true rank of the easy dataset is $R = 10$ (excluding correlations in the noise). The presence of outliers make the true rank N, although the effective rank (measured by its eigenvalue spectrum) is much closer to R. Figure 2 illustrates the 'corrupted' dataset.

3.8.1 Emulator performance

Figure 3 describes performance (as assessed by the measures from Section 3.6) over 20 replicates of the simulated datasets. Results are shown for the Mixture Model and Nearest Neighbour Emulators, choosing objects via the RMSE Loss, Furthest Distance Heuristic and Random Choice Heuristic. These are examined as a function of the number of iterations t and as a function of the dataset difficulty.

The emulator performance story is consistent across all choice functions. The difficulty of the problem is an important factor. Figure 3a-b shows that the simpler problems with easier cluster structure demonstrate the largest difference between methods. Nearest-neighbour emulation works well for small t, whilst the space is still being explored, but quickly asymptotes as the number of clusters is reached.

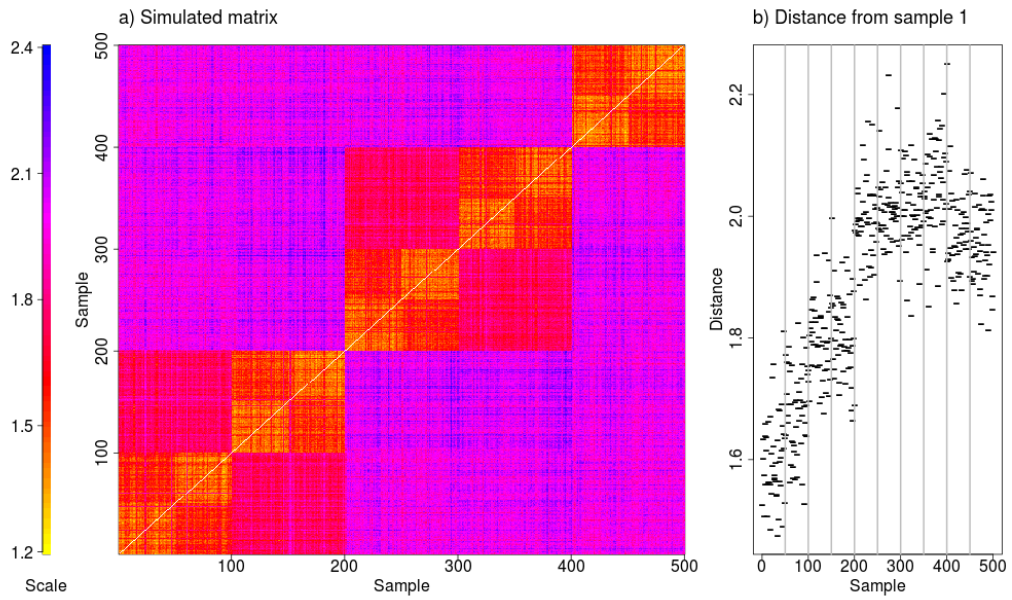

Figure 2: Simulated dataset, showing a) the whole 'corrupted' matrix as a heatmap, and b) the distribution of distances from an arbitrary row. Note that neighbouring clusters overlap strongly, and next-nearest clusters have overlapping distributions. Outliers appear as lower intensity rows and columns, and clusters contain 50 samples.

The Mixture model performs well overall and has very low asymptotic variance for this problem. Notably, direct matrix completion approach performs significantly worse than any approach we considered. The algorithm of Ma *et al.* (2011) achieved an RMSE $= 0.16$ using 25000 random matrix elements corresponding to $t = 50$. For this reason it is not considered further, although in the discussion we explain that the standard matrix completion problem is slightly misspecified for our problem.

The impact of the decision criterion on the whole matrix completion is dramatic as shown in Figure 3. For clustered datasets and early in the process, the furthest distance criterion performs very well. This success comes from first selecting objects from different clusters, but also selects outliers. For large numbers of iterations, this can be worse than choosing at random. The RMSE loss procedure, as might be expected, has the best overall performance and is robust to outliers and corrupted data. Similarly, the emulator variance estimate $\delta_i(t)$ depend on both the Choice function and emulator.

Figure 3c-d shows how well the emulator variance is captured. Unlike the predictive performance, this asymptotes to a finite, but small, value.

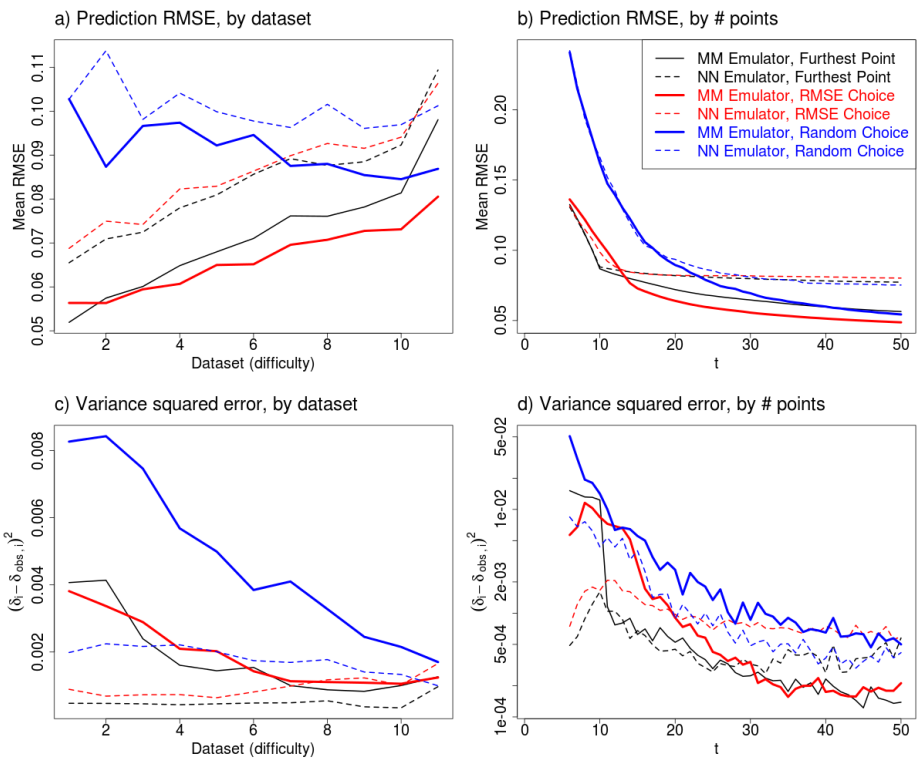

Figure 3: Mean emulation performance averaged over 20 simulated datasets. Left: performance as a function of the complexity of the dataset, averaged over iteration number (from 1 to 50), from 1='clustered' to 11='corrupted'. Right: Performance as a function t, the number of objects evaluated. a-b) shows the root-mean-square prediction error (RMSE) for the unobserved portion of the matrix. c-d) shows the emulator error when predicting the 'uncertainty' $\delta_i(t)$ in the estimate of the t-th iteration. Shown are both the Nearest-Neighbour and Mixture Model prediction methods, with objects selected using either the furthest distance heuristic, RMSE loss estimation, or random choice heuristic.

Throughout, using the RMSE choice criterion and hence the Emulator dominates Random Selection Heuristic (shown in thicker lines). It also dominates the Furthest Distance Heuristic except during the early phase ($t < 12$).

4 Computational scaling

We are interested in the rate of scaling up to logarithmic terms, which in practice are negligible compared to scaling constants. Fundamentally our approach allows

for linear scaling in N, provided that a) the matrix is of sufficiently low rank, and b) correct decision making for which elements to select requires only $O(1)$ calculation, compared to N. This is true for the Random Selection Heuristic and the Furthest Distance Heuristic. The methods using matrix completion, i.e. the maximum variance criterion and the root-mean-square-error criterion, scale significantly less well, though we restricted to $O(N)$ computation by estimating the loss on a $p \times p$ matrix.

The emulator used also has an important impact on running time. Calculating the nearest neighbour of an object to a set of t objects is an $O(t)$ operation, whereas inferring a mixture model using quadratic programming is only guaranteed to be polynomial Monteiro *et al.* (1990) in t. In practice it is $O(t^2)$ and hence of practical value. Therefore we suggest using this algorithm at the start of a computation, when making good decisions can strongly influence the decision making process, and is relatively efficient.

Algorithm	Order
Random Selection Heuristic	NT_{max}
Furthest Distance Heuristic	NT_{max}
Mixture Model, RMSE Loss	$NT_{max}^2p^2$
Nearest Neighbour, RMSE Loss	$NT_{max}p^2$

Table 2: Comparison of algorithmic scaling as a function of the number of objects N, the number to be evaluated T_{max} and the number p to be used in estimating the RMSE.

These arguments are summarised in Table 2, with Figure 4a-b confirming the results. We show the computational cost C_{choice} of using various decision models as a function of either the total number of objects N or the number to be evaluated T_{max}. We have constructed the code to make evaluation of S_{ij} approximately negligible so that the scaling of the decision framework dominates. Hence, random decisions take only the amount of time required to define the memory for the data, to copy it, and maintain state (all models share this cost).

The important criterion is the total compute $C_{choice} + C_{calculation}$, with the cost of evaluating the matrix elements being $C_{calculation} = LT_{max}N$. Hence is it vital to contrast the decision making cost to the calculation cost. Figure 4c-f show several scenarios of calculation cost whilst varying the problem difficulty. This leads to different optimal choices of decision strategy. Assuming a problem size of $N = 10000$, we consider two costs (measured in seconds) per similarity evaluation; $L = 0.01$ (Scenarios 1-2) and $L = 0.001$ (Scenarios 3-4). The first results in a final cost that is larger the most expensive (RMSE with a mixture model) decision cost by a factor of 5, whereas the second is within the range of computational budgets. By considering the RMSE curves from Figure 3, we can create curves comparing total computational cost to final RMSE, regardless of

T_{max}. By construction, this leads to a varying optimal decision; when L is large, it pays to use the best model for most values of computational cost. When L is small, it is optimal to either use a cheap heuristic (furthest neighbour) or random samples, depending on their performance in the problem.

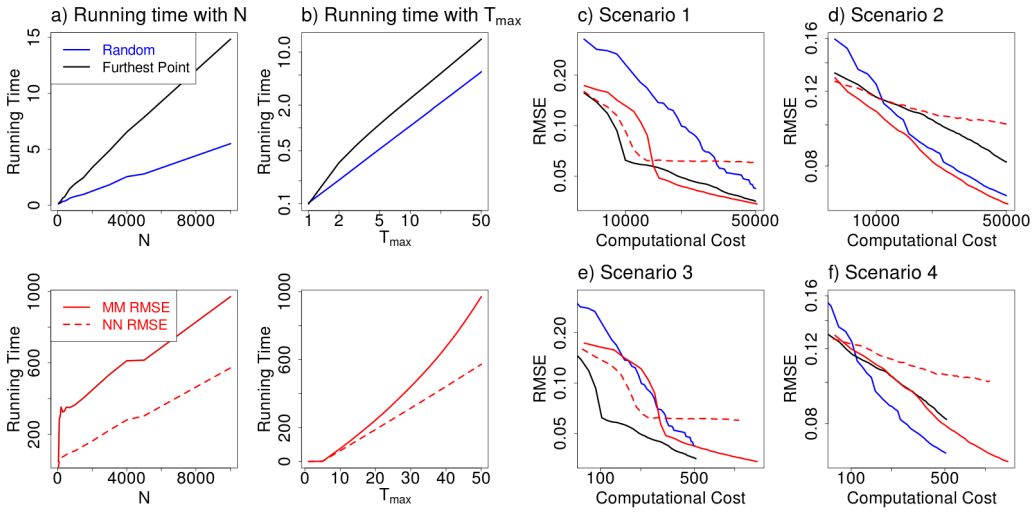

Figure 4: The scaling and performance considerations of computational cost. a) shows how the running time depends on the matrix size N ($T_{max} = 50$) for four configurations of our algorithm. These are the Random Choice Heuristic, the Furthest Distance Heuristic, and the RMSE Loss using either the Mixture Model or Nearest Neighbour emulators. b) shows the same results for T_{max} (at $N = 10000$). Note the different scales between the top and bottom plots. c-f show RMSE as a function of computational time (note the logarithmic x and y axis). This is under different hypothetical scenarios about computational cost and model efficiency (again assuming $N = 10000$). Scenarios 1 and 2 assume that $L = 0.1$ seconds per matrix element computation, whereas Scenarios 3 and 4 assume that $L = 0.001$. Scenarios 1 and 3 use the RMSE(t) curves from 'Dataset 1' (clustered) as shown in Figure 3, whereas Scenarios 2 and 4 use 'Dataset 11' (corrupted). Low RMSE for a given computational cost is desirable; the scenarios have been constructed to illustrate different optimal choices.

Figure 4a-b shows that our approach does indeed scale well, with our implementation in R handling large matrices ($N = 10000$ has 100 million elements) in completely acceptable times. The practical limitation is that we hope to make good decisions, which for large matrices would need heuristics that can be performed in $O(N)$ time or less. Performing careful calculations on an $O(1)$ fraction of the matrix will not be effective on general similarity matrices; we need a set of heuristics to pull out the most useful objects to consider. Prior information, in

the form of covariates or labels, could be very valuable for this. Xu *et al.* (2013) use such 'side information' in a standard matrix completion framework.

These computational scaling results do not include matrix completion of the final set of objects. This can be performed with any set of chosen objects using any available emulator. The cost of this is $O(N^2 T_{max})$ for nearest neighbour and $O(N^2 T_{max}^2)$ for the mixture model.

Figure 4c-f illustrate several scenarios for which the different computational costs L are accounted for. If L is moderate (Scenarios 1-2, costing $L = 0.1$ seconds per matrix element computation) then the RMSE Loss is best. This is universally true for difficult problems whilst the Furthest Distance Criterion is better for small budgets in the clustered dataset, when not much of the matrix is evaluated. Conversely, if L is very small (Scenarios 3-4 with $L = 0.001$ seconds), heuristics are necessary since computing matrix elements is as cheap as emulating them. In Scenario 3 with low L and clustered dataset, the Furthest Distance Heuristic is unbeatable, whereas in Scenario 4 the corrupted dataset means that the Random Choice Heuristic is better for moderate computational budget.

5 Genetics example

An important application in statistical genetics is the understanding of the relationship between individuals, and the overall structure of that relationship. The process generating genetics data is very well understood - a generative model called the ancestral recombination graph (ARG, Griffiths (1981); Hudson (1983)) operates within populations. Between populations, migration occurs at varying rates over time Lawson (2014). Because of this, model-based inference is strongly preferred by this community. However, the true underlying genetic model is prohibitively computationally costly for all but completely unrealistically small and simple samples. Therefore much effort goes into the development of approximation procedures that capture the key features of the ARG for a manageable computational cost.

We are interested in a population assignment problem Lawson *et al.* (2012), for which the likelihood for population assignment can be shown to approximately follow a similarity matrix. It is therefore of great importance to approximate this matrix. Fundamentally, the difficulty is that computing the similarity matrix cost $O(N^2 L)$ with L taking values up to around 20 million for whole genome resequencing data. The number of individuals sampled N is also growing increasingly large. $N > 1000$ is becoming standard (the first being the 1000 Genome project Consortium *et al.* (2012), now with 2500 comparable whole genomes), whilst $N = 10000$ lower density samples are in existence (e.g. Ripke *et al.* (2013)), and larger datasets are on their way. Computing a row of similarities for this matrix takes days of computation time, and although parallelization is possible,

most genetics departments do not have access to the 1000+ core compute farms needed to avoid serial computation. Thus, there is a strong motivation to reduce the computation of these methods, whilst the fundamental computation cannot be changed. Since N^2 and L are of the same order, we will use an $O(NL + N^2)$ implementation of the algorithm.

We apply the method to the well understood Human Genome Diversity Panel dataset formed of $N = 938$ individuals from across the globe and $L \approx 2 \times 800,000$ single nucleotide polymorphism (SNP) data points per individual. This is large enough to be difficult to work with but well within the scope of institutional clusters, and was explored in the original Lawson *et al.* (2012) paper. Our goal here is to recreate the essential features of the matrix at a fraction of the computational cost. The true rank of this matrix is at least 226 as this is the number of populations identified using full model based inference. This problem contains all of the difficulties of the 'hard' simulated data and several more: a) it is not symmetric; b) some individuals are much less similar to any individuals than others; c) it is high rank; d) clusters are of all sizes.

We first converted the genetic similarity to a distance by taking the negative log and adding the new minimum value. We then applied the RMSE Choice criterion with the Mixture Model emulator using cross-validation to estimate $\delta_i(t)$. Up to 100 samples were used to recover the matrix. The performance is shown in Figure 5, compared to both the Random Choice Heuristic and a new 'Prior Heuristic'. This exploits prior information about populations using self-declared ethnicities. To do this we sampled one individual from each population, and then a second individual from each population, etc, with both populations and individuals within those populations being defined at random.

6 Discussion

Massive data sets present new challenges for statistical reasoning. The sheer volume of data prevents the routine deployment of even straightforward statistical techniques. Thus a different paradigm is required. Our proposed framework attacks this problem using statistical reasoning to construct a sequential data selection paradigm. At core, this framework acknowledges that it is impossible to complete the statistical evaluation on all data, and instead seeks to find the most useful subset in a sequential manner. This is most suited to problems in which correlation prevents the use of simple sampling, such as the AR model example in Section 3.1.3, or the central example of similarity matrices.

The framework combines elements of active learning, modelling, and sequential decision making. At a high-level, the framework features a statistical emulator (for cheaply estimating unknown data values), a loss, and a choice operator (for selecting among candidates to determine the next computation). We have provided

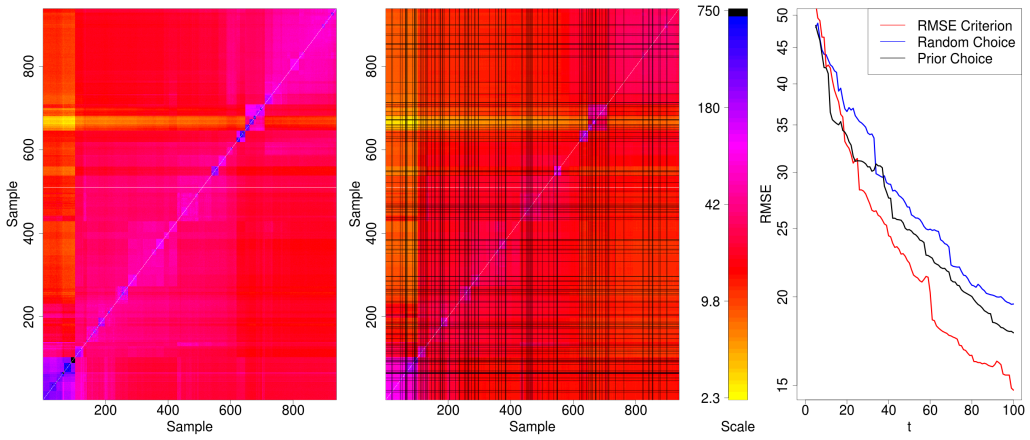

Figure 5: Genetics example. Left: the raw data ordered as Figure 4A and S14 of Lawson *et al.* (2012). Centre: The recovered matrix after 100 iterations, showing the individuals used for the recovery. Right: the root-mean-square-error of the matrix recovery over time, for the RMSE choice criterion, compared to random choice and exploiting the prior information of the 53 population labels by ensuring that each label is sampled in a balanced order. Lower is better.

a number of choices for each, which have different capabilities and computational costs. In the case of similarity matrices, understanding the latent subspace is usually advantageous. However, with massive data this task is difficulty. Hence we have introduced Data Directional Scaling a method of determining important directions in similarity space itself. This approach allows a rapid characterisation of the main directions without appeal to a costly latent variable model.

Whilst we explored many advantages of data directional scaling, the key disadvantage is that the dimension of the space grows with the number of samples. It would be fairly simple to extend our decision framework to 'discard' directions of the data that are not helpful for prediction. A more sophisticated approach is to perform online clustering to create pseudo-directions via K-means Hartigan and Wong (1979) or the more rapid K-medians Jain and Dubes (1988). This step brings with it a potential predictive performance drop, since we provide the emulator with less information, and would also need to solve the serious problem of estimating K. Further, we cannot compute $\mathcal{S}(i, j)$ for these pseudo-objects although we can still use them in the emulator. Hence we have omitted it from this work, although addressing these concerns is vital to replace the computational scaling factor T_{max} by K in our emulator steps.

We have provided examples in which our approach is to be preferred over matrix completion approaches. This preference is not merely empirical. Matrix

completion may still require an intractable computation at large enough data scale, hence precluding its use. Worse, the model assumptions of matrix completion may not match the characteristics of many similarity matrices, in which cluster structure is often expected. Assessing empirical performance is complicated since we have to handle issues of computational cost, choice function etc. For the type of challenging problem with which we are concerned, where L and N are large, it is important to compare against the benchmark of random selection. In such cases, empirical results show that the framework is far superior to random selection.

Further development of this framework will follow in two directions. On the abstract side, we seek some theoretical guarantees on the behaviour of the framework. Moreover, refinements of the details of the framework, emulator and choice operator, should lead to performance enhancements. On the computational side, extension of the approach to parallel and cloud systems is of great interest. The big data we are interested in will usually reside on a distributed storing system, such as HADOOP.

References

Bastos, L. S. and A. O'Hagan, 2009 Diagnostics for Gaussian process emulators. Technometrics *51*(4)**:** 425–438.

Bayarri, M., J. O. Berger, M. C. Kennedy, A. Kottas, R. Paulo, J. Sacks, J. A. Cafeo, C.-H. Lin, and J. Tu, 2009 Predicting vehicle crashworthiness: Validation of computer models for functional and hierarchical data. Journal of the American Statistical Association *104*(487).

Bentley, J. L., 1975 Multidimensional binary search trees used for associative searching. Communications of the ACM *18*(9)**:** 509–517.

Berchtold, S., D. A. Keim, and H.-P. Kriegel, 2001 The X-tree: An index structure for high-dimensional data. Readings in multimedia computing and networking**:** 451.

Berger, J. O., 1985 *Statistical decision theory and Bayesian analysis*. Springer.

Bhojanapalli, S. and P. Jain, 2014, February)Universal Matrix Completion. ArXiv e-prints.

Bijak, J., J. Hilton, E. Silverman, and V. D. Cao, 2013 Reforging the Wedding Ring: Exploring a Semi-Artificial Model of Population for the United Kingdom with Gaussian process Emulators. Demographic Research **29**.

Biswas, P., T.-C. Lian, T.-C. Wang, and Y. Ye, 2006 Semidefinite programming based algoirithms for sensor network localization. ACM Trans. Sens. Netw. *2*(2)**:** 188–220.

Busby, D. and M. Feraille, 2008 Adaptive design of experiments for calibration of complex simulators–An application to uncertainty quantification of a mature oil field. In *Journal of Physics: Conference Series*, Volume 135, pp. 012026. IOP Publishing.

Candes, E. J. and Y. Plan, 2010 Matrix completion with noise. Proceedings of the IEEE *98*(6)**:** 925–936.

Cesa-Bianchi, N. and G. Lugosi, 2006 *Prediction, Learning and Games*. Cambridge University Press.

Chi, E. C., H. Zhou, G. K. Chen, D. O. Del Vecchyo, and K. Lange, 2013 Genotype imputation via matrix completion. Genome Research *23*(3)**:** 509–518.

Cohn, D. A., Z. Ghahramani, and M. I. Jordan, 1996 Active Learning with Statistical Models. Journal of Artificial Intelligence Research **4:** 129–145.

Consortium, . G. P. and others, 2012 An integrated map of genetic variation from 1,092 human genomes. Nature *491*(7422)**:** 56–65.

Finkel, R. A. and J. L. Bentley, 1974 Quad trees a data structure for retrieval on composite keys. Acta informatica *4*(1)**:** 1–9.

Gamerman, D., 1997 *Markov Chain Monte Carlo: Stochastic simulation for Bayesian inference*. 2-6 Boundary Row, London, UK. SE1 8HN: Chapman and Hall.

Gerlee, P., L. Schmidt, N. Monsefi, T. Kling, and R. J. et al., 2013 Searching for Synergies: Matrix Algebraic Approaches for Efficient Pair Screening. PLoS ONE *8*(7).

Gramacy, R. B. and D. W. Apley, 2013 Local Gaussian process approximation for large computer experiments. arXiv preprint arXiv:1303.0383.

Griffiths, R. C., 1981 Neutral two-locus multiple allele models with recombination. Theor. Popul. Biol. **19:** 169–186.

Guttman, A., 1984 R-Trees: a dynamic index structure for spatial searching. In *Proceedings of the 1984 ACM SIGMOD international conference on Management of data - SIGMOD '84*, pp. 47–.

Hamilton, J. D., 1994 *Time series analysis*, Volume 2. Princeton university press Princeton.

Hao, N., L. Horesh, and M. Kilmer, 2014 Nuclear Norm Optimization and Its Application to Observation Model Specification. In A. Y. Carmi, L. Mihaylova, and S. J. Godsill (Eds.), *Compressed Sensing & Sparse Filtering*, Signals and Communication Technology, pp. 95–122. Springer Berlin Heidelberg.

Hartigan, J. A. and M. A. Wong, 1979 A k-means clustering algorithm. Applied Statistics **28:** 100–108.

Howell, D. C., 2007 The treatment of missing data. The Sage handbook of social science methodology**:** 208–224.

Hudson, R. R., 1983 Properties of the neutral model with intragenic recombination. Theor.Pop.Biol. **23:** 213–201.

Jain, A. K. and R. C. Dubes, 1988 *Algorithms for clustering data*. Prentice-Hall, Inc.

Jones, R. H., 1980 Maximum likelihood fitting of ARMA models to time series with missing observations. Technometrics *22*(3)**:** 389–395.

Kalman, R. E., 1960 A new approach to linear filtering and prediction problems. Journal of basic Engineering *82*(1)**:** 35–45.

Kennedy, M. C., C. W. Anderson, S. Conti, and A. OHagan, 2006 Case studies in Gaussian process modelling of computer codes. Reliability Engineering & System Safety *91*(10)**:** 1301–1309.

Keshavan, R. H., A. Montanari, and O. Sewoong, 2010 Matrix completion from a few entries. IEEE Transactions on Information Theory *56*(6)**:** 2980–2998.

Lawson, C. L. and R. J. Hanson, 1974 *Solving least squares problems*, Volume 161. SIAM.

Lawson, D. J., 2014 Populations in statistical genetics modelling and inference. In U. . W. Kreager, Capelli (Ed.), *Population in the Human Sciences; Concepts, Methods, Evidence*. OUP.

Lawson, D. J. and D. Falush, 2012 Population identification using genetic data. Annual Review of Genomics and Human Genetics **13:** 337–361.

Lawson, D. J., G. Hellenthal, S. Myers, and D. Falush, 2012 Inference of population structure using dense haplotype data. PLoS Genetics **8:** e100245.

Liao, Y., 2013 Learning to predict end-to-end network performance. Ph. D. thesis, Department of Electronics XXX!

Lindley, D. V., 1961 Dynamic programming and decision theory. Applied Statistics**:** 39–51.

Ma, S., D. Goldfarb, and L. Chen, 2011 Fixed point and Bregman iterative methods for matrix rank minimization. Mathematical Programming *128*(1-2)**:** 321–353.

Mazumder, R., T. Hastie, and R. Tibshirani, 2010 Spectral regularization algorithms for learning large incomplete matrices. The Journal of Machine Learning Research **11:** 2287–2322.

Monteiro, R. D., I. Adler, and M. G. Resende, 1990 A polynomial-time primal-dual affine scaling algorithm for linear and convex quadratic programming and its power series extension. Mathematics of Operations Research *15*(2)**:** 191–214.

Ricci, F., L. Rokach, B. Shapira, and P. Cantor (Eds.), 2011 *Recommender Systems Handbook*. Springer.

Ripke, S., N. R. Wray, C. M. Lewis, S. P. Hamilton, M. M. Weissman, G. Breen, E. M. Byrne, D. H. Blackwood, D. I. Boomsma, S. Cichon, and others, 2013 A mega-analysis of genome-wide association studies for major depressive disorder. Molecular psychiatry *18*(4)**:** 497–511.

Rougier, J., D. M. Sexton, J. M. Murphy, and D. Stainforth, 2009 Analyzing the Climate Sensitivity of the HadSM3 Climate Model Using Ensembles from Different but Related Experiments. Journal of Climate *22*(13).

Santner, T. J., B. J. Williams, and W. I. Notz, 2003 *The design and analysis of computer experiments*. Springer.

Seo, S., M. Wallat, T. Graepel, and K. Obermayer, 2000 Gaussian process regression: Active data selection and test point rejection. In *Mustererkennung 2000*, pp. 27–34. Springer.

Settles, B., 2010 Active learning literature survey. University of Wisconsin, Madison **52:** 55–66.

Simmonds, R., R. Bradford, and B. Unger, 2000 Applying parallel discrete event simulation to network emulation. In *Proceedings of the fourteenth workshop on Parallel and distributed simulation*, pp. 15–22. IEEE Computer Society.

Šochman, J. and J. Matas, 2009 Learning fast emulators of binary decision processes. International Journal of Computer Vision *83*(2)**:** 149–163.

Taghizadeh, M. J., 2014 Theoretical Analysis of Euclidean Distance Matrix Completion for Ad hoc Microphone Array Calibration. In *Idiap Technical Report*.

Thompson, J. D., D. G. Higgins, and T. J. Gibson, 1994 CLUSTAL W: improving the sensitivity of progressive multiple sequence alignment through sequence weighting, position-specific gap penalties and weight matrix choice. Nucleic acids research *22*(22)**:** 4673–4680.

Turlach, B. A. and A. Weingessel, 2013 quadprog: Functions to solve Quadratic Programming Problems. R package version 1.5-5.

Xu, M., R. Jin, and Z.-H. Zhou, 2013 Speedup Matrix Completion with Side Information: Application to Multi-Label Learning. In C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger (Eds.), *Advances in Neural Information Processing Systems 26*, pp. 2301–2309.

Young, F. W. and R. M. Hamer, 1987 *Multidimensional Scaling: History, Theory and Applications*. New York: Erlbaum.